arXiv:1511.08225v1 [physics.comp-ph] 25 Nov 2015

Speeding Up Simulations By Slowing Down Particles: Speed-Limited Particle-In-Cell Simulation☆

Gregory R. Werner[a,*], John R. Cary[a,b]

[a] *Center for Integrated Plasma Studies, University of Colorado, Boulder, CO 80309*
[b] *Tech-X Corporation, Boulder, CO 80303*

Abstract

Particle-in-cell (PIC) simulation is often impractical for the same reason that it is powerful: it includes too much physics. Sometimes the mere ability to simulate physics on small length or time scales requires those scales to be resolved (by the cell size and timestep) to avoid instability, even when the effects at those scales contribute negligibly to the phenomenon motivating the simulation. For example, a timestep larger than the inverse plasma frequency will often result in unphysical growth of plasma oscillations, even in simulations where plasma oscillations should not arise at all. Larger timesteps are possible in simulations based on reduced physics models, such as MHD or gyrokinetics, or in simulations with implicit time-advances. A new method, speed-limited PIC (SLPIC) simulation, allows larger timesteps without reduced physics and with an explicit time-advance. The SLPIC method slows down fast particles while still accurately representing the particle distribution. SLPIC is valid when fields and distribution functions change slowly compared with the desired timestep; SLPIC is useful when that timestep is much larger than that allowed by standard PIC (and when alternative approximations do not include enough physics). Speed-limited PIC can be implemented with relatively localized modifications of a standard PIC code.

Keywords: simulation, FDTD, PIC, finite difference time domain, particle-in-cell, speed-limited particle-in-cell

1. Introduction

Particle-in-cell (PIC) simulation of plasmas is a powerful technique for studying plasma phenomena, in large part because it can include all of the classical "first principles" physics—i.e., the Lorentz force and Maxwell's equations (with the latter sometimes helpfully simplified, e.g., to Poisson's equation) [1, 2]. To accomplish this, PIC simulations track sample "macroparticles," which follow trajectories of real particles and represent (parts of) particle distribution functions. PIC simulation is of course still an

☆GW was partially supported by the Institute for Modeling Plasma, Atmospheres, and Cosmic Dust (IMPACT) of NASA's Solar System Virtual Institute (SSERVI).
*Corresponding author: Greg.Werner@colorado.edu

approximation of the underlying physics; however, in the limit of infinite grid resolution, infinitesimal timestep, and small (hence numerous) macroparticles, PIC simulation includes all the physics of Maxwell's equations and the Lorentz force law.

The very strength of PIC is sometimes its biggest drawback, because simulations are often limited by what one *can* simulate rather than what one *wants* to simulate. For example, because electrostatic PIC simulation *can* capture the phenomenon of plasma oscillation, simulations experience instability unless the timestep Δt is small enough to resolve the plasma frequency [3]. With a timestep determined by the plasma frequency, PIC simulation is often too costly (computationally) to simulate phenomena that operate on much longer timescales than plasma oscillations.

A new PIC-based simulation technique—speed-limited PIC (SLPIC)—modifies the PIC method to slow down fast phenomena, enabling larger timesteps while retaining the same underlying physics on slow timescales. SLPIC has the potential to speed up simulation when there exist fast phenomena that are numerically troublesome, but physically unimportant. For such simulations SLPIC improves upon PIC by increasing the amount of approximation; however, the strength of approximation can be continuously varied until SLPIC is identical to PIC, making it possible—without leaving the SLPIC framework or sacrificing efficiency—to verify whether faster phenomena actually have negligible effect

There are two main classes of methods developed to suppress numerical problems related to irrelevant physics: reduced-physics methods that neglect (or integrate over) irrelevant physics, and methods that use a special time-advance to avoid growth of unresolved plasma modes. The first class can be very useful when a reduced physics model (e.g., magnetohydrodynamics or gyrokinetics) is known and is known to be applicable; however, extending such models to include more physics tends to be difficult or impossible. The second class (e.g., fully implicit time-advances) can be difficult to implement and computationally expensive, often involving iterative nonlinear solves. SLPIC does not require a reduced physics model, and can be explicit; moreover, it can be implemented as a minor modification to an existing PIC code. SLPIC is not a universal replacement for other methods, but it does extend the feasible range of PIC simulation to longer timescales. SLPIC has the potential to be the fastest method in a niche where kinetic simulation is required and distribution functions change relatively slowly compared to the plasma frequency.

Fluid-based simulations are important examples of the first class; they use reduced physics models to avoid numerical problems with irrelevant phenomena. For example, magnetohydrodynamics (MHD) simulations lack the ability to simulate plasma oscillations, so they avoid any timestep limitation related to the plasma frequency. Another fluid approach—using Boltzmann electrons—treats electrons as a fluid in thermal equilibrium so that the electron density is given by $n_0 \exp(e\Phi/kT)$, where Φ is the local electrostatic potential. The use of Boltzmann electrons allows PIC (hence fully kinetic) treatment of ions and ion dynamics while neglecting most electron dynamics (see, e.g., [3, 4]). SLPIC may be useful in circumstances where the Boltzmann electron approach is almost sufficient, but more physics is required. SLPIC allows the electron distribution to relax to the ion distribution (as Boltzmann electrons relax to the potential determined by the ions), but SLPIC can evolve arbitrary electron distributions. For example, we expect SLPIC to speed simulation of phenomena such as collisionless sheaths (in which the electron distribution is not Maxwellian) and electron Landau damping of ion-acoustic

waves.

Some (kinetic) PIC-based approaches, such as gyrokinetics (cf. [5]), modify the particle equations of motion to integrate over small length and time scales. SLPIC is somewhat similar in that it also averages over fast phenomena by modifying particle equations of motion, but differs from gyrokinetics because the SLPIC approximation does not affect the physics (i.e., does not use a reduced physics model) at slow timescales. We note that SLPIC could also be applied to a reduced physics model such as gyrokinetics; however, SLPIC's advantage may rely on the simplicity and power obtained from using the complete (classical) physics.

In principle, SLPIC is a method for modifying any PIC algorithm (e.g., whether standard Lorentz-force or gyrokinetic) by slowing down fast particles. In SLPIC, slow particles behave just as in normal PIC, while fast particles are "speed-limited" and locally follow the same phase-space trajectories as real particles but in slow motion. In the limit of sufficiently slowly-varying fields, SLPIC particles follow the same phase-space trajectories as real particles (over finite times, not just locally), but at different speeds. It is practical for the speed-limit v_0 (separating fast and slow particles) to depend on the timestep, $v_0 \propto 1/\Delta t$, and thus SLPIC introduces a continuously-variable approximation (in addition to the PIC approximation) that depends on the relative scales of the desired timestep and the temporal variation of particle distribution functions. As the timestep approaches the (stable) PIC timestep, SLPIC simulation becomes identical (in function and computational cost) to PIC simulation.

Although SLPIC can be (and is probably most effective when) used with an explicit time-advance, it shares some similarities with fully implicit PIC methods (see, e.g., [6, 7]) that allow larger timesteps without reduced physics or instability. Naturally and unavoidably, such implicit methods inaccurately simulate phenomena that are poorly resolved by the timestep; unlike explicit methods, however, implicit methods may remain stable even when not resolving irrelevant high-frequency phenomena. As with implicit methods, the approximation in SLPIC is continuously adjustable through the choice of timestep; the choice of timestep in SLPIC as in implicit PIC is a choice to simulate faster phenomena inaccurately—a choice that is justified when those phenomena are unimportant.

SLPIC has several possible advantages over implicit PIC. First, SLPIC can be explicit, hence faster than implicit PIC (for the same timestep). Second, SLPIC is very similar to standard PIC (i.e., much simpler than implicit PIC), requiring little modification to algorithms except those governing individual macroparticle trajectories. Third, SLPIC handles the problem of particles crossing too many grid cells within a timestep: crossing multiple cells poses a challenge for parallel computation, which is necessary for most large plasma simulations.

SLPIC offers a method that (unlike MHD or Boltzmann electrons) is essentially similar to PIC and does not require any reduced physics model. An existing PIC code can be modified to support SLPIC with relatively little trouble; for example, field solvers remain completely unchanged. The main difference is that (fast) particles move in slow motion. The most prevalent case where SLPIC can speed simulation is perhaps when the electron distribution relaxes (to a quasi-equilibrium) on ion time scales.

SLPIC involves explicit, local modifications of a standard PIC code; while extra computation is required, that extra computation is local (involving only an individual particle and its equation of motion) and predictable/consistent (i.e., not iterative, or

requiring new solvers that might affect scaling with problem size or number of parallel processors).

SLPIC is a very new and promising simulation technique. We introduce and justify the fundamental approach in the following section; in subsequent sections we show calculation of the plasma frequency in SLPIC and demonstrate the effectiveness of SLPIC for collisionless sheath simulation.

2. SLPIC

The goal of *kinetic* simulation is the self-consistent evolution of the particle distribution function $f(x, v, t)$, which in a collisionless plasma satisfies the Vlasov equation,

$$\partial_t f(x, v, t) + \nabla_x \cdot [v f(x, v, t)] + \nabla_v \cdot [a(x, v, t) f(x, v, t)] = 0, \tag{1}$$

where $a(x, v, t)$ is acceleration due to whatever forces act on a particle located at x with velocity v at time t. (Although we present this analysis with zero on the right-hand side above, a non-zero value, e.g., due to collisions, would not alter the SLPIC technique.) For a Hamiltonian (phase-space-preserving) system, $\nabla_x \cdot v + \nabla_v \cdot a = 0$, leading to the familiar form,

$$\partial_t f(x, v, t) + v \cdot \nabla_x f(x, v, t) + a \cdot \nabla_v f(x, v, t) = 0. \tag{2}$$

Equation (1) can be solved by the method of characteristics, since it describes an element of the phase-space distribution advecting through phase space with the local velocity and acceleration. Thus, it has a solution as a sum over particles (or trajectories),

$$f(x, v, t) = \sum_p w_p \delta(x - x_p(t)) \delta(v - v_p(t)), \tag{3}$$

where x_p and v_p are particle trajectories (and w_p is a weight representing the number of real particles embodied in macroparticle p)—i.e., they satisfy

$$\dot{x}_p = v_p, \tag{4}$$

$$\dot{v}_p = a(x_p, v_p, t). \tag{5}$$

[One can verify by direct substitution that, with these equations of motion, Eq. (3) is a solution of Eq. (1).] This is the basis of PIC simulation [1, 2]; PIC methods will not be reviewed here, except to say that they in essence broaden the delta-function $\delta(x - x_p(t))$ so that particle charges and currents can be transferred to a grid for calculation of fields (the scatter operation), while fields on a grid are interpolated to particles' positions (the gather operation) to yield the forces on the particles.

While this approach can in principle simulate all the fundamental (classical) physics of plasmas, the separation of scales—especially between electron and ion motion—often renders practical simulation impossible with current resources. Two important (and related) problems make simulation slow: (1) the timestep must generally be smaller than the grid-cell-crossing time of the fastest particle, and (2) the timestep must be smaller than the inverse plasma frequency. The first reason is related to accuracy (as well as practicality for parallel computing): since fields cannot vary on length scales

smaller than a grid cell, a particle traversing less than one cell-length experiences only small changes in fields. The second condition is crucial to avoid catastrophic numerical instability: with a timestep greater than $2/\omega_e$, where ω_e is the plasma frequency, the standard and simple leap-frog integration scheme is unstable, with numerical solutions that grow exponentially.

If one also wishes to avoid the grid instability (unphysical heating that increases the Debye length until it is resolved by the grid), one must usually choose a grid cell that resolves the electron Debye length, λ_e, in the case of a stationary, thermal plasma. In this case, the two criteria above are identical within a factor of order unity.

For cases where time scales of interest are long compared with the plasma period, the above conditions on the timestep are prohibitive. E.g., for steady state plasma sheaths, the relaxation time is of the order of ion time scales, which are longer by $\sim \sqrt{m_i/m_e}$. For such simulations we would like to choose a timestep of the order of the ion plasma period, but we are prevented from doing so by numerical instability. In such cases, the electrons are effectively in equilibrium with the electrostatic field. Critically, they are in a kinetic equilibrium, but a nontrivial one, since in a collisionless sheath problem, electrons that flow into the sheath with enough energy to overcome the sheath potential do not reflect back into the plasma or equilibrate, and so a Boltzmann dependence $n_e \sim \exp(e\Phi/kT_e)$ is not accurate.

To address such situations we propose using "speed-limited" electrons, which reduce the scale-separation when the dynamics of interest take place over times that are long compared with the inverse plasma frequency or the cell-crossing time of the fastest particles. To do this we limit the *speed* with which simulated electrons travel through the simulation to some maximum v_0, but preserve the correct *direction* of travel, ensuring that a speed-limited electron follows the same path as a real electron (but at a slower speed). We will show that this approach allows larger timesteps, hence faster simulation, while accurately capturing the physics of longer time scales—including kinetic effects of electrons on those timescales.

To use this method, we will simulate a distribution $g(x, v, t)$ of speed-limited (macro)particles, defined through the speed-limiting factor $\beta(x, v, t)$ [we will see later in Eq. (11) that the speed-limited velocity is β times the real velocity, $v_s = \beta v$]:

$$f(x, v, t) = \beta(x, v, t)g(x, v, t). \tag{6}$$

Whereas PIC evolves $f(x, v, t)$, SLPIC uses PIC methods to evolve $g(x, v, t)$, which can be trivially converted to $f(x, v, t)$ at any time. Inserting this into Eq. (1) yields

$$\partial_t[\beta g(x, v, t)] + \nabla_x \cdot [\beta v g(x, v, t)] + \nabla_v \cdot [\beta a(x, v, t)g(x, v, t)] = 0, \tag{7}$$

which we rewrite in the form,

$$\partial_t g(x, v, t) + \nabla_x \cdot [\beta v g(x, v, t)] + \nabla_v \cdot [\beta a(x, v, t)g(x, v, t)] = \partial_t[(1 - \beta)g(x, v, t)], \tag{8}$$

The approximation that makes SLPIC useful is the neglect of the right-hand side above; in the limit that the right-hand side vanishes, SLPIC is as accurate as PIC. Therefore, SLPIC achieves full-PIC accuracy in the limit of static perturbations (steady-state), where all the time derivatives vanish—and in particular the right side of Eq. (8) vanishes. Also, when $\beta = 1$, the right-hand side vanishes and again SLPIC achieves exactly the same accuracy as standard PIC.

Effective use of SLPIC involves choosing $\beta(x,v,t)$ such that for $|v| \ll v_0$, $\beta = 1$ and so slow particles are simulated just as accurately as in full PIC, while particles with $|v| > v_0$ are simulated under the approximation that the time-derivative on the right-hand side is negligible. The latter is a good approximation when the distribution function changes sufficiently slowly.

In more detail, we consider (not necessarily small) perturbations or plasma modes with characteristic (phase) velocities v_ϕ. Particles moving much faster than such perturbations equilibrate with them rapidly. However, particles with velocities of order v_ϕ are not in equilibrium with the perturbations. Indeed, such particles can lead to Landau damping, extracting energy from the perturbations. Hence for such particles all temporal derivatives in the above equation must be kept. However, if we set β to be nearly unity for velocities less than v_0 where $v_0 \gtrsim v_\phi$, then the right side of Eq. (8) vanishes for low-velocity particles also. Thus, it is a uniform approximation to set the right-hand side of Eq. (8) to zero:

$$\partial_t g(x,v,t) + \nabla_x \cdot [\beta v g(x,v,t)] + \nabla_v \cdot [\beta a(x,v,t) g(x,v,t)] = 0, \tag{9}$$

since it is valid for high velocities by accurately giving their equilibrium, and it is valid for low velocities because $\beta \approx 1$.

With this approximation, we can evolve $g(x,v,t)$ according to

$$g(x,v,t) \quad = \quad \sum_p w_p(t) \delta(x - x_p(t)) \delta(v - v_p(t)) \tag{10}$$

by the method of characteristics, just as Eq. (1), except that to satisfy Eq. (9) the macroparticles are slow (or speed-limited), following trajectories $x_p(t)$, $v_p(t)$ satisfying the equation of motion (hereafter we omit the particle index p)

$$\dot{x} \quad = \quad v_s = \beta(x,v,t) v, \tag{11}$$

$$\dot{v} \quad = \quad a_s = \beta(x,v,t) a(x,v,t). \tag{12}$$

SLPIC evolves $g(x,v,t)$ in time—and this (because of the SLPIC approximation) is computationally faster than evolving $f(x,v,t)$ because the macroparticles move more slowly, allowing larger timesteps. With a known $\beta(x,v,t)$, one can convert g to f as desired.

In some sense, β acts like a macroparticle weight; i.e., a macroparticle that is used to evolve g is subsequently "weighted" by $\beta(x,v,t)$ to compute f. With this view, f is the sum of macroparticles following trajectories given by Eq. (11) with weights changing according to $(d/dt)\beta$ evaluated along each particle's trajectory.

There are many possibilities for the speed-limiting function $\beta(x,v,t)$. For SLPIC, β needs to limit the speed of macroparticles to some value v_0. For large velocities, we must have $\beta \approx v_0/|v|$ (to limit $|\dot{x}|$ to v_0), and for small velocities, $\beta \approx 1$ [so the right-hand side of Eq. (8) vanishes]. Examples are

$$\beta = \frac{v_0}{\sqrt{v^2 + v_0^2}} \qquad \text{or} \qquad \beta = \Theta(v_0 - |v|) + \frac{v_0}{|v|} \Theta(|v| - v_0) \tag{13}$$

where Θ is the Heaviside step function, but there are many other options. It is possible even for β to have spatial dependence, e.g., for a spatially varying grid, or time-dependence, e.g, to adjust the severity of approximation mid-simulation. However, we

will denote the speed limit by v_0, keeping in mind that it may vary with position (and possibly time).

An important aspect of this method is that, since the right-hand sides of Eqs. (4) and (11) differ by a scalar factor β, speed-limited particles [representing $g(x, v, t)$] follow the same phase-space trajectories as real particles [representing $f(x, v, t)$], except at a slower speed. Therefore, even though fast electrons (with $v \gg v_0$) move unphysically slowly (with $\dot{x} \sim v_0$) under SLPIC, they have the correct response as long as the fields evolve slowly. As a speed-limited particle (with $v \gg v_0$) accelerates, its actual velocity v increases (so it gains energy), but $\dot{x}$ remains near v_0; to compensate, its weight must decrease. E.g., in a steady-state streaming fluid, an increase in velocity results in a decrease in density; in SLPIC, the macroparticle speed doesn't change much, so the macroparticle density doesn't change much, and the macroparticle weight decreases to reflect the real decrease in density.

With particles limited to speeds below v_0, one may choose the timestep $\Delta t \lesssim \Delta x/v_0$, where Δx is the cell size, so that particles will not cross more than one cell per timestep. This allows an increase in timestep by a factor $v_{\max}/v_0 \sim v_e/v_0$, where $v_{\max}$ is the maximum particle speed and v_e is the thermal velocity. Like methods involving Boltzmann electrons, this method is useful when the electron distribution is quasi steady-state on the timescales of ion motion. However, unlike Boltzmann electron methods, this method simulates an arbitrary electron distribution and has an "adjustable" approximation, which allows the simulation to change continuously into a full PIC simulation by increasing v_0 just above the maximum particle speed.

Despite some mathematical resemblance, SLPIC is not a δf-PIC method, nor is it equivalent to lowering the speed of light to v_0. Whereas δf methods evolve a perturbed distribution $\delta f(x, v, t)$ on top of a given (usually equilibrium) distribution $f_0(x, v)$ (see, e.g., [8]), SLPIC evolves the entire distribution function $f(x, v, t)$ and does not require that the solution be a small perturbation of some known equilibrium. And while lowering the speed of light to v_0 would certainly impose a speed limit, particles would not follow the same trajectories as if the speed of light had not been lowered.

3. Plasma oscillations for speed-limited electrons

As we noted above, using speed-limited electrons allows us to relax the cell-crossing timestep restriction, as electrons move more slowly through the simulation. It turns out that speed-limiting of electrons also lowers the electron plasma frequency, which relaxes the other condition that required a small timestep. Here we show that the plasma frequency for speed-limited electrons is reduced by $\sim v_0/v_e$ (again, allowing the timestep to be increased by a factor of v_e/v_0, where v_e is the electron thermal velocity).

To compute the plasma frequency in the SLPIC system, we consider 1D wave-like perturbations $\exp[i(kx - \omega t)]$ from a zero-field, uniform, steady-state distribution g_0 with $\beta = \beta(v)$ independent of space and time. Denoting unperturbed quantities with subscript 0, and first-order with subscript 1, the first-order solution to Eq. (9) for the speed-limited distribution function is

$$-i(\omega - k\beta v)\tilde{g_1} = -\tilde{a_1}\partial_v[\beta g_0(v)], \tag{14}$$

where a_1 is the acceleration due to the first-order electric field, and over-tildes indicate amplitudes of oscillation. I.e.,

$$g_1(x,v,t) = \tilde{g_1}(v)\exp(ikx - i\omega t). \tag{15}$$

From this we find the solution for the density perturbation:

$$\begin{aligned} \tilde{n_1} &= \int dv\, \tilde{f_1} = \int dv\, \beta \tilde{g_1} \;=\; -i\tilde{a_1}\int dv\, \frac{\beta}{\omega - k\beta v}\partial_v[\beta g_0] \\ &= i\tilde{a_1}\int dv\, \beta g_0 \partial_v \frac{\beta}{\omega - k\beta v} \;=\; i\tilde{a_1}\int dv\, f_0 \partial_v \frac{\beta}{\omega - k\beta v} \\ &= i\tilde{a_1} n_0 \left\langle \partial_v \frac{\beta}{\omega - k\beta v} \right\rangle, \end{aligned} \tag{16}$$

where the angled brackets denote the average over the velocity distribution function f_0. Inserting this into Gauss's law [$\partial_x E_1 = (-e)n_1/\epsilon_0$ or $ik\tilde{a_1} = \tilde{n_1}e^2/(m\epsilon_0)$] we find

$$1 \;=\; \frac{\omega_p^2}{k}\left\langle \partial_v \frac{\beta}{\omega - k\beta v} \right\rangle \;=\; \frac{\omega_p^2}{\omega^2}\left\langle \frac{\beta^2 + (\partial_v\beta)\omega/k}{(1 - k\beta v/\omega)^2} \right\rangle, \tag{17}$$

where $\omega_p^2 = e^2 n_0/(m\epsilon_0)$. The plasma frequency is found by looking at the long wavelength limit $k \to 0$. When $\beta = 1$ (hence $\partial_v\beta = 0$), we recover the standard result: $\omega = \omega_p$.

When $\partial_v\beta \neq 0$, then for isotropic f_0 and β [i.e., $f_0(-v) = f_0(v)$ and $\beta(v) = \beta(-v)$], $\int dv f_0(v)\partial_v\beta(v) = 0$, and expansion of the denominator to first order yields

$$\omega^2 = \omega_p^2 \left\langle \beta^2 + 2\beta\beta' v \right\rangle. \tag{18}$$

When the limiting velocity v_0 is much less than the electron thermal velocity v_e, the speed-limiting function is $\beta \sim v_0/|v|$ and the speed-limited plasma frequency ω_s is

$$\omega_s \sim \omega_p \frac{v_0}{v_e}. \tag{19}$$

I.e., the effective plasma frequency is reduced by nearly the same fraction by which a typical particle's speed is limited, v_0/v_e.

By speed-limiting simulated particles to v_0, i.e., reducing the speed of typical electrons by v_0/v_e, we reduce the simulated plasma frequency by a similar factor. Naturally, plasma oscillations are not accurately simulated, but that is an advantage because this method is appropriate only for cases where the important dynamics are much slower than plasma oscillations.

3.1. SLPIC for a 1D plasma sheath

To prove its basic usefulness and accuracy, we implemented speed-limited PIC simulation in 1D and applied it to collisionless sheath simulation. For sheath simulation, the Debye length must be resolved and (in normal PIC simulation) the timestep must be shorter than the time for the fastest electron to cross a cell, $\Delta t_{\rm pic} \lesssim \Delta x/(3v_e)$. On the other hand, the sheath will approach a steady state on the time scale with which slow ions cross the entire simulation. For a simulation that is $10\lambda_D$ with $\Delta x = \lambda_D/10$, the

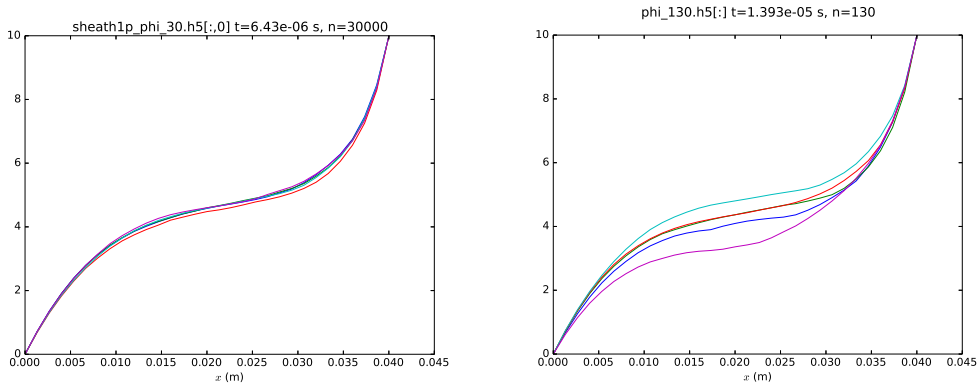

Figure 1: The sheath potential at 5 different times, separated by 1 μs intervals, just after reaching a steady state: (left) normal PIC simulation, which reached a steady-state after 10^4 steps; and (right) SLPIC simulation, which reached a steady-state after 90 steps. The SLPIC simulation is noisier, for reasons yet to be understood (the time-ordering of the curves is: blue, green, red, cyan, magenta).

(hydrogen) ion crossing time can be on the order of 10^4 timesteps (and 6 times worse for, say, argon ions).

On the other hand, with speed-limited PIC simulation, we can greatly increase the timestep, allowing much faster simulation. For a very simple test case of a sheath simulation, we chose $T_e = 2.2$eV and $T_i = 0.5$eV, and an extremely low $v_0 = v_e/100 = 0.9v_i$ (so even many ions are speed-limited). While speed-limiting typical ions is probably not desirable because it is not likely to reduce the time to solution, it should not harm the steady state in principle, and so we chose this v_0 to demonstrate this principle. With speed-limiting, we ran the simulation with a timestep 520 times normal; $\Delta t_{\text{slpic}} = 520\Delta t_{\text{pic}}$. Our python toy code running the speed-limited simulation reached a steady state after approximately 90 steps (or 10μs physical time), while the normal PIC simulation (Vorpal, [9]) reached a steady state after about 10^4 steps (or 2μs physical time). Because our extremely low v_0 was less than the ion thermal velocity, it's not surprising that the speed-limited simulation evolved more slowly in physical time.

The resulting sheath potentials are shown in Fig. 1. The similarity in sheath potentials—with the 500$\times$ increase in time-step—is very promising.

4. Summary

Speed-limited particle-in-cell (SLPIC) simulation is a new technique that allows kinetic PIC simulation with a larger timestep, by limiting the maximum speed of particles (which otherwise follow the proper trajectories). The choice of speed limit controls the strength of the approximation introduced by SLPIC: as the speed-limit increases beyond the speed of the fastest particle, SLPIC becomes identical to PIC. Lowering the speed limit also lowers the plasma frequency, which allows the timestep to be increased without courting the instability that results in standard PIC from the failure to resolve the plasma frequency; it also prevents particles from traveling too far during one (increased) timestep.

SLPIC overcomes a common limitation of explicit PIC simulation: even when the phenomena of interest (hence fields and distribution functions) do not involve plasma oscillations and change slowly compared with the plasma frequency, the PIC timestep must be small enough to resolve the plasma frequency to avoid instability. SLPIC allows the timestep to be determined by the phenomena of interest, rather than by the irrelevant and much faster plasma oscillations.

SLPIC promises to be useful in cases where particle distribution functions change slowly compared with the timestep required by PIC for stability (and also for preventing particles from crossing too many cells in one timestep). For example, SLPIC may be especially applicable to cases where the electron distribution relaxes on ion time scales.

In a typical SLPIC simulation with a cell size approximately equal to a Debye length, one might imagine setting the speed limit some factor below typical electron speeds, but above physically-important velocities such as the phase velocity of a plasma mode being simulated. The timestep could then be increased by the same factor, speeding up computation by roughly that factor, reduced slightly because the SLPIC particle-advance requires more operations than standard PIC. The extra operations are local (to a single particle) and explicit, so they will take advantage of memory cache and may be especially amenable to hardware acceleration; and they will not affect overall scaling with problem size or computational resources.

The speed limit can in principle vary in both space and time; this offers intriguing possibilities of increasing the speed-limit (and decreasing the timestep) mid-simulation to increase and/or verify the accuracy of approximation.

Modification of a PIC code to implement SLPIC is expected to be localized to the integration of individual particle trajectories, and consideration of particle weight (e.g., in charge deposition). Other aspects of PIC, such as collisions, field solvers, boundary conditions, etc., should carry over to SLPIC with little or no change. SLPIC therefore promises a particularly flexible and powerful approach to increase the timestep of PIC simulations.

The basic SLPIC method has been tested in a 1D simulation of a steady-state plasma sheath, successfully yielding an accurate potential profile despite a timestep 500 times larger than would be used in standard PIC simulation.

References

[1] C. K. Birdsall, A. B. Langdon, Plasma physics via computer simulation, CRC Press, 2004.

[2] R. W. Hockney, J. W. Eastwood, Computer simulation using particles, CRC Press, 1988.

[3] R. Mason, Computer simulation of ion-acoustic shocks. The diaphragm problem, Phys. Fluids 14 (9) (1971) 1943.

[4] K. Cartwright, J. Verboncoeur, C. Birdsall, Nonlinear hybrid Boltzmann–particle-in-cell acceleration algorithm, Phys. Plasmas 7 (8) (2000) 3252.

[5] W. Lee, Gyrokinetic approach in particle simulation, Phys. Fluids 26 (2) (1983) 556.

[6] S. Markidis, G. Lapenta, The energy conserving particle-in-cell method, J. Comput. Phys. 230 (18) (2011) 7037.

[7] L. Chacón, G. Chen, D. Barnes, A charge-and energy-conserving implicit, electrostatic particle-in-cell algorithm on mapped computational meshes, J. Comput. Phys. 233 (2013) 1.

[8] S. Parker, W. Lee, A fully nonlinear characteristic method for gyrokinetic simulation, Physics of Fluids B: Plasma Physics (1989-1993) 5 (1) (1993) 77.

[9] C. Nieter, J. R. Cary, VORPAL: a versatile plasma simulation code, J. Comput. Phys. 196 (2) (2004) 448.